arXiv:1602.01347v1 [stat.OT] 3 Feb 2016

Determining optimal factors for chemical synthesis of pharmaceutical products using experimental data

Antony M. Overstall
School of Mathematics and Statistics,
University of Glasgow,
Glasgow G12 8QW UK
(Antony.Overstall@glasgow.ac.uk)

David C. Woods
Southampton Statistical Sciences Research Institute,
University of Southampton,
Southampton SO17 1BJ UK
(D.Woods@southampton.ac.uk)

Kieran J. Martin
Roche

Abstract

In a chemical synthesis process to manufacture a pharmaceutical product, an initial set of substances evolve according to chemical reactions, under certain process conditions, into a series of new substances. One of these substances is a target pharmaceutical product and two are unwanted by-products. The aim is to determine the factors (process conditions and amounts of initial substances) that maximise the probability of the amounts of pharmaceutical product and by-products being greater and less than certain levels, respectively. The relationship between the factors and amounts of substances of interest is theoretically described by the intractable solution to a system of ordinary differential equations incorporating temperature dependence. This relationship is refined using observations from an experiment by the fitting of a statistical model. Predictions from this model are used to evaluate the probability of satisfying the constraints and this probability is maximised over the set of factors to provide optimal manufacturing conditions for the pharmaceutical product. To accelerate the model-fitting and the evaluation of the probability of satisfying the constraints, the numerical solution to the differential equations is approximated using statistical emulators.

Keywords: approximate coordinate exchange; experiments; multivariate Gaussian process; prediction; Riemann manifold Langevin Metropolis-Hastings

1 Introduction

Pharmaceutical products are often manufactured by chemical synthesis. This is where an initial set of chemical substances evolve in time into a substance or substances of interest via a series of chemical reactions. One (or more) of these substances of interest will be the target pharmaceutical product while the others will be unwanted or harmful *by-product* substances. The chemical synthesis process may be hypothesised to depend on various controllable factors (and time). The aim of the chemical engineers will be to manipulate the controllable factors such that the amounts of interest satisfy certain constraints. For example, that the amounts of the substances that make up the pharmaceutical product are above some specified level whilst the amounts of unwanted substances are correspondingly below a certain level.

The exact relationship between the controllable factors and the synthesis process is unknown. However, in some cases, an experiment can be performed where the amounts of the substances of interest are measured for several different specifications of the controllable factors. The observed data can then be used to inform the relationship between controllable factors and the synthesis process.

In this paper, we consider an example of such an exercise. There are three substances of interest: one of which constitutes the pharmaceutical product and the other two are unwanted by-products. The dynamics behind the synthesis process are thought to be described by a system of non-linear ordinary differential equations (ODEs), the solution to which forms a physical model for the amounts of the three substances at a given time. This model depends on a set of controllable factors (including time) and unknown parameters. A small experiment has been conducted where the amounts of the substances of interest have been observed for different controllable factors at certain time points. The aim is to use this observed data to estimate the unknown parameters and then to use these estimates to maximise (over the set of controllable factors and time) the probability of the predicted amounts satisfying some inequality constraints. The approach used here will be Bayesian giving a coherent method of propagating uncertainty on the unknown parameters (given the observed data) through to the predicted amounts.

There are several challenges to this approach. First, the solution to the system of ODEs is analytically intractable and so requires a computationally expensive numerical solution (e.g. Iserles, 2009) to be employed. This will affect both the estimation of the unknown parameters through the evaluation of their posterior distribution *and* in making predictions for new controllable factors. To alleviate this problem we make use of (multivariate) Gaussian process emulators to accelerate both the Markov chain Monte Carlo (MCMC) algorithm employed and the prediction of future amounts. Second, some of the observations fall below a minimum threshold, χ, under which the actual amount cannot be measured reliably. Thus we have left-censored observations. Instead of assuming that the true observations are just equal to zero or χ, we assume that the true observations are unknown, subject to being in the interval $(0, \chi)$, in a natural Bayesian formulation. Third, we know that the physical model derived from the solution to the ODEs only provides an approximation to the true process. We account for this mismatch by introducing an explicit model discrepancy term (Kennedy and O'Hagan, 2001).

The paper is organised as follows. In Section 2 we describe the experiment, and the physical and statistical models for the observations in more detail. In Section 3 we describe the methodology including how the Gaussian process emulators are constructed. In Section 4 we present the results of applying the methodology to the real application before concluding with a discussion in Section 5.

2 Background

2.1 Chemical Reactions

The chemical synthesis process involves nine substances labelled A to I. The substances of interest are E, F and H, where F constitutes the pharmaceutical product and E and H are the by-product substances. Let $A(t)$ denote the amount (in mols) of substance A at time t (in seconds) with similar notation for the remaining substances. At time $t = 0$, initial amounts of A, D and E, denoted by $A_0 = A(0)$, $D_0 = D(0)$ and $E_0 = E(0)$, respectively, are introduced into the laboratory apparatus at a specified temperature (λ in °C) and volume (V, in l). Let $\mathbf{x} = (A_0, D_0, E_0, \lambda, V)$ be an $F \times 1$ vector ($F = 5$) and assume that $\mathbf{x} \in \mathcal{X} = \prod_{f=1}^{F} [x_{f,min}, x_{f,max}]$. The values of $x_{f,min}$ and $x_{f,max}$ are given in Table 1.

The chemical synthesis process is governed by the following series of chemical reactions:

$$\left.\begin{array}{lcl} \mathrm{A} & \xrightarrow{k_1} & 2\mathrm{B} + \mathrm{C} \\ \mathrm{B} + \mathrm{D} + \mathrm{E} & \xleftrightarrow{k_2} & \mathrm{B} + \mathrm{G} + \mathrm{F} \\ \mathrm{F} + \mathrm{B} & \xrightarrow{k_3} & \mathrm{H} + \mathrm{I} \end{array}\right\} \tag{1}$$

where k_1, k_2, and k_3 denote the chemical reaction rates. The second line denotes a reversible reaction and thus k_2 is split into k_2^- and k_2^+, for the forwards and backwards reactions, respectively.

Table 1: Lower and upper limits of $\mathbf{x}$

Definition	Name	Lower limit, $x_{f,min}$	Upper limit, $x_{f,max}$	Units
Initial amount of A	$x_1 = A_0$	22.5	40	mols
Initial amount of D	$x_2 = D_0$	91.4	91.59	mols
Initial amount of E	$x_3 = E_0$	26.42	26.47	mols
Temperature	$x_4 = \lambda$	25	40	°C
Volume	$x_5 = V$	31.28	32.56	l

Let $[A](t) = A(t)/V$ be the concentration (in mol/l) at time t of A, let $\dot{A}(t) = \mathrm{d}[A](t)/\mathrm{d}t$ be the corresponding time derivative, and assume similar definitions for the remaining substances; B to I. The chemical reactions given by (1) lead to the following system of non-linear ODEs (dropping the dependence on t),

$$\left.\begin{array}{rcl} \dot{A} &=& -k_1[A] \\ \dot{B} &=& 2k_1[A] - k_3[F][B] \\ \dot{C} &=& k_1[A] \\ \dot{D} &=& -k_2^-[B][D][E] + k_2^+[B][G][F] \\ \dot{E} &=& -k_2^-[B][D][E] + k_2^+[B][G][F] \\ \dot{F} &=& k_2^-[B][D][E] - k_2^+[B][G][F] - k_3[F][B] \\ \dot{G} &=& k_2^-[B][D][E] - k_2^+[B][G][F] \\ \dot{H} &=& k_3[F][B] \\ \dot{I} &=& k_3[F][B] \end{array}\right\} \tag{2}$$

for $t \in \mathcal{T} = [0, 3000]$s. The initial concentrations of all nine substances are zero, except for A, D and E which are $[A](0) = A_0/V$, $[D](0) = D_0/V$ and $[E](0) = E_0/V$, respectively.

For a given chemical reaction rates, the solution, $([A](t), \ldots, [I](t))$, to the system given by (2) provides the concentrations of the nine substances for given controllable factors. The amounts of the nine substances is then given by $(A(t), \ldots, I(t)) = V \times ([A](t), \ldots, [I](t))$. However the solution is analytically intractable. We employ a computationally expensive numerical method to find an approximate solution, namely that given by the `deSolve()` (Soetaert et al., 2010) function in the statistical software package `R` (R Core Team, 2015).

The question is what values of the initial amounts of A, D and E, the temperature, the volume and time should be chosen such that

$$\left.\begin{array}{rcl} E(t) &<& 3, \\ F(t) &>& 20, \\ H(t) &<& 3? \end{array}\right\} \tag{3}$$

The chemical engineers have assumed that $k_1 = 10000$ but the remaining reaction rates, k_2^-, k_2^+ and k_3, are unknown. Temperature dependence is incorporated into the system of ODEs via the Arrehnius equation (see, for example, Laidler 1984). The general form for a reaction rate is

$$k_i = k_i^{(r)} \exp\left(\frac{E_i}{G}\left(\frac{1}{\lambda} - \frac{1}{\lambda^{(r)}}\right)\right),$$

where $\lambda^{(r)}$ is the reference temperature (here 25°C), $k_i^{(r)} > 0$ is the reaction rate at the reference temperature, $E_i > 0$ is the activation energy and $G = 8.31445$ Jmol^{-1}K^{-1} is the gas constant. By definition, the activation energies for the reversible reactions are equal, i.e. $E_2^+ = E_2^- = E_2$. This means there are $p = 5$ unknown physical parameters, $\boldsymbol{\gamma} = \left(k_2^{(r)-}, k_2^{(r)+}, k_3^{(r)}, E_2, E_3\right)$. To help in the implementation of the methodology described in Section 3, we actually operate on the logarithm scale, i.e. consider $\theta_i = \log \gamma_i$, for $i = 1, \ldots, p$, and let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_p) \in \Theta = \mathbb{R}^p$.

2.2 Experiment and Physical Model

We estimate the values of the unknown physical parameters $\boldsymbol{\theta}$ using observations from a small, non-optimally designed experiment. The chemical reactions given by (1) are observed for $I' = 6$ combinations (or treatments) of the controllable factors $\mathbf{x}$. These treatments are given in Table 2. For run $i = 1, \ldots, I'$, we observe the amounts of E, F and H at a series of n_i times $t_{i1}, \ldots, t_{in_i}$. The times range from 0.5 to 2902 seconds, with n_i ranging from 17 to 20 time points per run. In total, there are $n = \sum_{i=1}^{I'} n_i = 109$ observations of the amounts over all I' runs. Note how run 6 is a repetition of run 5, thus there are $I = I' - 1 = 5$ unique treatments. Runs 5 and 6 have observations taken at 18 different time points. However the last two time points for run 5 are at 1620 and 1825 seconds, whereas the last two time points for run 6 are taken four seconds later. Therefore there are $m = 93$ unique combinations of controllable factors and time points.Let $\left(\mathbf{x}'_j, t_j\right)$ be the values of the controllable factors for each $(j = 1, \ldots, m)$ of these unique combinations.

Table 2: Six treatments of the controllable factors, $\mathbf{x}$, for the experiment

		Initial amount (in mols)				
Run	Symbol	A_0	D_0	E_0	Temperature (λ,°C)	Volume (V, l)
1	$\mathbf{x}_1$	22.50	91.59	26.47	25.00	31.31
2	$\mathbf{x}_2$	45.00	91.59	26.47	25.00	32.56
3	$\mathbf{x}_3$	22.50	91.50	26.45	40.00	31.28
4	$\mathbf{x}_4$	45.00	91.50	26.45	40.00	32.53
5	$\mathbf{x}_5$	33.75	91.40	26.42	32.50	31.88
6	$\mathbf{x}_6$	33.75	91.40	26.42	32.50	31.88

Let K be the number of substances of interest, i.e. $K = 3$. Let $\boldsymbol{\mu}\left(\boldsymbol{\theta}; \mathbf{x}, t\right) = \log\left(E(t), F(t), H(t)\right)$ be the $K \times 1$ vector giving the log of the theoretical amount of the three substances of interest (E, F and H) based on the numerical solution to the system of ODEs, at $\mathbf{x}$, time t and parameters $\boldsymbol{\theta}$.

Let $\mathbf{V}_i$ be the $n_i \times K$ matrix containing the observed amounts of E, F and H for the ith run, for $i = 1, \ldots, I'$. The jth row of $\mathbf{V}_i$, denoted by $\mathbf{v}_{ij}$ is a $K \times 1$ vector containing the observed amounts for the jth time point of the ith run, for $i = 1, \ldots, I'$ and $j = 1, \ldots, n_i$. Let $\mathbf{V}$ be the $n \times K$ matrix formed by stacking the matrices $\mathbf{V}_1, \ldots, \mathbf{V}_{I'}$, i.e.

$$\mathbf{V} = \begin{pmatrix} \mathbf{V}_1 \\ \vdots \\ \mathbf{V}_{I'} \end{pmatrix}.$$

The experimental apparatus cannot detect substance amounts lower than $\chi = 0.01$ mols. Therefore, we have censored values when the observed amount is less than χ. In total, there are zero censored observations in the first column of $\mathbf{V}$, and 5 and 45 in the second and third columns, respectively. We take account of the censoring in our statistical modelling strategy described in the Section 2.3.

We further define the $n \times K$ matrix $\mathbf{Y}$ by taking the natural logarithm of each element of $\mathbf{V}$.

2.3 Statistical Model

We assume that

$$\mathbf{Y} = \mathbf{G}\left(\mathbf{M}(\boldsymbol{\theta}) + \mathbf{D}\right) + \mathbf{E}. \tag{4}$$

In (4), $\mathbf{M}(\boldsymbol{\theta})$ is the $m \times K$ matrix given by

$$\mathbf{M}(\boldsymbol{\theta}) = \begin{pmatrix} \mathbf{M}_1(\boldsymbol{\theta}) \\ \vdots \\ \mathbf{M}_I(\boldsymbol{\theta}) \end{pmatrix},$$

where $\mathbf{M}_i(\boldsymbol{\theta})$, for $i = 1, \ldots, I-1$, is the $n_i \times k$ matrix with jth row given by $\mu(\boldsymbol{\theta}; \mathbf{x}_i, t_{ij})$, and $\mathbf{M}_I(\boldsymbol{\theta})$ is the $(n_I + 2) \times k$ matrix where the jth row is given by $\mu(\boldsymbol{\theta}; \mathbf{x}_I, t_{Ij})$, for $j = 1, \ldots, n_I$, and the $(n_I + 1)$th and $(n_I + 2)$th rows are given by $\mu(\boldsymbol{\theta}; \mathbf{x}_I, 1624)$ and $\mu(\boldsymbol{\theta}; \mathbf{x}_I, 1829)$, respectively.

Furthermore, $\mathbf{D}$ is the $m \times K$ matrix of model discrepancy errors given by

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}_1 \\ \vdots \\ \mathbf{D}_I \end{pmatrix},$$

which represents the discrepancy between the mechanistic model and reality (Kennedy and O'Hagan, 2001). Here $\mathbf{G}$ is the $n \times m$ matrix of zeros and ones which identifies the rows of $\mathbf{M} + \mathbf{D}$ which refer to the corresponding rows of $\mathbf{Y}$. Finally, $\mathbf{E}$ is the $n \times K$ matrix of observational errors.

The distribution of $\mathbf{E}$ is assumed to be

$$\mathbf{E} \sim \mathrm{MN}(\mathbf{0}, \boldsymbol{\Omega}, \mathbf{T}),$$

where MN denotes the matrix normal distribution (e.g. Overstall and Woods, 2016), $\boldsymbol{\Omega}$ is an unknown $K \times K$ column covariance matrix, and $\mathbf{T}$ is an $n \times n$ row correlation matrix. The matrix $\mathbf{T}$ is block-diagonal, given by $\mathbf{T} = \mathrm{diag}\{\mathbf{T}_1, \ldots, \mathbf{T}_{I'}\}$, where the jlth element of $\mathbf{T}_i$ is

$$T_{i,jl} = \exp\left(-\rho\left(t_{ij} - t_{il}\right)^2\right),$$

with the time correlation parameter $\rho > 0$ is unknown. This covariance structure for $\mathbf{E}$ assumes that the correlation between observations from the same run is dependent on the difference in time, but observations from different runs are independent.

The distribution of $\mathbf{D}$ is assumed to be

$$\mathbf{D} \sim \mathrm{MN}(\mathbf{0}, \boldsymbol{\Sigma}, \mathbf{S}), \tag{5}$$

where $\boldsymbol{\Sigma}$ is an unknown $K \times K$ covariance matrix and $\mathbf{S}$ is an $m \times m$ correlation matrix. The jlth element of $\mathbf{S}$ is given by

$$S_{jl} = \exp\left(-\psi_1 \sum_{f=1}^{F} \left(d_{jp} - d_{lp}\right)^2 - \psi_2 \left(t_j - t_l\right)^2\right), \tag{6}$$

for $j, l = 1, \ldots, m$ and where the model discrepancy correlation parameters $\psi_1 > 0$ and $\psi_2 > 0$ are unknown. In (6), $d_{jf} = \left(x'_{jf} - x_{f,min}\right) / \left(x_{f,max} - x_{f,min}\right)$ is the scaled (divided by range) value of the fth controllable factor, for $f = 1, \ldots, F$ and $j = 1, \ldots, m$, where x'_{jf} is the fth element of $\mathbf{x}'_j$. This correlation structure takes account of how the model discrepancy will be dependent on the "distance" between the controllable factors and time points, and implies that the model discrepancy is not independent between different runs.

We complete the model specification with the prior distribution for the model parameters; $\boldsymbol{\theta}$, $\boldsymbol{\Omega}$, $\boldsymbol{\Sigma}$, ρ, ψ_1 and ψ_2. The chemical engineers have some limited prior knowledge on the location of $\boldsymbol{\theta}$ and this is approximately represented by the following prior distribution:

$$\log \boldsymbol{\theta} \sim \mathrm{N}(\boldsymbol{\mu}, \boldsymbol{\Delta}),$$

where $\boldsymbol{\mu} = (-13, -13, -13, 10, 10)$ and $\boldsymbol{\Delta} = \mathrm{diag}(\boldsymbol{\kappa})$ with $\boldsymbol{\kappa} = (2, 2, 2, 3, 3)$. There is no prior information available for the remaining parameters so we need to specify prior distributions that contribute negligible information. The variance matrices $\boldsymbol{\Omega}$ and $\boldsymbol{\Sigma}$ are both assumed to have inverse-Wishart prior distributions with $\nu = 4$ degrees of freedom and an identity scale matrix (Gelman et al., 2014, Chapter 3). The correlation parameters ρ, ψ_1 and ψ_2 are given independent exponential prior distributions with mean equal to one following the arguments of, for example, Overstall and Woods (2016).

Let $\mathbf{y} = \text{vec}(\mathbf{Y})$, $\mathbf{m}(\boldsymbol{\theta}) = \text{vec}(\mathbf{M}(\boldsymbol{\theta}))$, $\mathbf{d} = \text{vec}(\mathbf{D})$ and $\mathbf{e} = \text{vec}(\mathbf{E})$, then an equivalent model specification to (4) is given by

$$\mathbf{y} = \mathbf{H}\left(\mathbf{m}(\boldsymbol{\theta}) + \mathbf{d}\right) + \mathbf{e}, \tag{7}$$

where $\mathbf{H} = \mathbf{I}_k \otimes \mathbf{G}$,

$$\mathbf{e} \sim \text{N}\left(\mathbf{0}, \boldsymbol{\Omega} \otimes \mathbf{T}\right), \tag{8}$$
$$\mathbf{d} \sim \text{N}\left(\mathbf{0}, \boldsymbol{\Sigma} \otimes \mathbf{S}\right), \tag{9}$$

and $\otimes$ denotes the Kronecker product.

Let $\mathbf{y}_S$ and $\mathbf{y}_C$ denote the elements of $\mathbf{y}$ which are fully observed (i.e. greater than χ) and censored (i.e. less than χ), respectively. Therefore, we need to evaluate the posterior distribution (conditional on $\mathbf{y}_S$) of the model parameters and censored observations. This distribution is given by

$$\pi(\boldsymbol{\theta}, \boldsymbol{\Omega}, \rho, \mathbf{d}, \boldsymbol{\Sigma}, \boldsymbol{\psi}, \mathbf{y}_C|\mathbf{y}_S) \propto \pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d}, \boldsymbol{\Omega}, \rho)\pi(\mathbf{d}|\boldsymbol{\Sigma}, \boldsymbol{\psi})\pi(\boldsymbol{\theta})\pi(\boldsymbol{\Omega})\pi(\rho)\pi(\boldsymbol{\Sigma})\pi(\boldsymbol{\psi}), \tag{10}$$

where $\pi(\mathbf{y}|\boldsymbol{\theta}, \mathbf{d}, \boldsymbol{\Omega}, \rho)$ is the complete data likelihood given by (7) and (8), and $\pi(\mathbf{d}|\boldsymbol{\Sigma}, \boldsymbol{\psi})$ is the distribution of the vectorised model discrepancy (9).

2.4 Prediction

Let $\mathbf{y}_0$ be a $K \times 1$ vector denoting the predicted log amounts of E, F and H for arbitary controllable factors $\mathbf{x} \in \mathcal{X}$ and $t \in \mathcal{T}$. The posterior predictive distribution of $\mathbf{y}_0$ is given by

$$\pi(\mathbf{y}_0|\mathbf{y}_S) = \int \pi(\mathbf{y}_0|\boldsymbol{\theta}, \mathbf{D}, \boldsymbol{\psi}, \boldsymbol{\Omega}, \boldsymbol{\Sigma})\pi(\boldsymbol{\theta}, \mathbf{D}, \boldsymbol{\psi}, \boldsymbol{\Omega}, \boldsymbol{\Sigma}|\mathbf{y}_S)\text{d}\boldsymbol{\theta}\text{d}\mathbf{D}\text{d}\boldsymbol{\psi}\text{d}\boldsymbol{\Omega}\text{d}\boldsymbol{\Sigma}, \tag{11}$$

where

$$\mathbf{y}_0|\boldsymbol{\theta}, \mathbf{D}, \boldsymbol{\psi}, \boldsymbol{\Omega}, \boldsymbol{\Sigma} \sim \text{N}\left(\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t) + \mathbf{D}^T\mathbf{S}^{-1}\mathbf{s}, \left(1 - \mathbf{s}^T\mathbf{S}^{-1}\mathbf{s}\right)\boldsymbol{\Sigma} + \boldsymbol{\Omega}\right),$$

and $\mathbf{s}$ is an $m \times 1$ vector with jth element given by

$$s_j = \exp\left(-\psi_1 \sum_{f=1}^{F} \left(d_{jf} - d_f\right)^2 - \psi_2 \left(t_j - t_f\right)^2\right),$$

with $d_f = \left(x_f - x_{f,min}\right) / \left(x_{f,max} - x_{f,min}\right)$ being the scaled value of each element of $\mathbf{x}$, for $f = 1, \ldots, F$.

The probability of E, F and H satisfying the constraints is given by

$$\text{P}(\mathbf{y}_0 \in \mathcal{Y}|\mathbf{y}_S) = \int_{\mathcal{Y}} \pi(\mathbf{y}_0|\mathbf{y}_S)\text{d}\mathbf{y}_0,$$

where $\mathcal{Y} = \left\{\boldsymbol{\eta} \in \mathbb{R}^K | \eta_1 < \log 3, \eta_2 > \log 20, \eta_3 < \log 3\right\}$.

2.5 General Approach

The posterior and posterior predictive distributions given by (10) and (11), respectively, will be analytically intractable.

The general approach will be to use the following two phases. The **Sampling Phase** involves generating a sample from the posterior distribution using MCMC methods. The **Prediction Phase** uses this MCMC sample to generate a sample from the posterior predictive distribution and thus estimate the probability $\text{P}(\mathbf{y}_0 \in \mathcal{Y}|\mathbf{y}_S)$ by the proportion of sampled values satisfying the constraints.

We describe the computational strategy for achieving this in Section 3.

3 Methodology

3.1 Gibbs sampling and parallel tempering

To generate a sample from the posterior distribution of model parameters and censored observations we use Gibbs sampling in conjunction with parallel tempering.

To improve the efficiency of the Gibbs sampler, we employ hierarchical centring (Papaspiliopoulos et al., 2003) and reparameterise model (7) as

$$\mathbf{y} = \mathbf{H}\mathbf{d}^* + \mathbf{e}, \tag{12}$$

where

$$\mathbf{d}^* \sim \mathrm{N}\left(\mathbf{m}(\boldsymbol{\theta}), \boldsymbol{\Sigma} \otimes \mathbf{S}\right). \tag{13}$$

Sampled values of $\mathbf{d}^*$ can easily be transformed to values of $\mathbf{d}$ using $\mathbf{d} = \mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})$. The posterior distribution of model parameters and censored observations is now given by

$$\pi(\boldsymbol{\theta}, \boldsymbol{\Omega}, \rho, \mathbf{d}^*, \boldsymbol{\Sigma}, \boldsymbol{\psi}, \mathbf{y}_C | \mathbf{y}_S) \propto \pi(\mathbf{y}|\mathbf{d}^*, \boldsymbol{\Omega}, \rho)\pi(\mathbf{d}^*|\boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\psi})\pi(\boldsymbol{\theta})\pi(\boldsymbol{\Omega})\pi(\rho)\pi(\boldsymbol{\Sigma})\pi(\boldsymbol{\psi}),$$

where $\pi(\mathbf{y}|\mathbf{d}^*, \boldsymbol{\Omega}, \rho)$ is the complete-data likelihood given by (12) and $\pi(\mathbf{d}^*|\boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\psi})$ is given by (13). The full conditional distributions of $\mathbf{d}^*$, $\boldsymbol{\Omega}$ and $\boldsymbol{\Sigma}$ are of known form and are given in A. The full conditional distribution of $\mathbf{y}_C$ is a truncated multivariate normal distribution (see A) and there exist efficient methods for generating values from such distributions (Geweke, 1991).

For the remaining parameters $\boldsymbol{\theta}$, ρ and $\boldsymbol{\psi}$, the full conditional distributions are not of known form so a Metropolis-within-Gibbs step is employed instead. Specifically, we exploit the recent development of the Riemann manifold Langevin Metropolis-Hastings (RMLMH) algorithm (Girolami and Calderhead, 2011) to complete this step. This method uses derivative information on the posterior surface to construct an efficient proposal distribution. The algorithmic details are briefly described in B and the components necessary for its application to the full conditional distributions of $\boldsymbol{\theta}$, ρ and $\boldsymbol{\psi}$ are given in C.

Parallel tempering (see, for example, Gelman et al., 2014, pages 299-300) is used to improve the efficiency of sampling from a potentially multi-modal posterior distribution. It involves setting up R distributions defined by a series of increasing temperatures where the lowest temperature distribution corresponds to the posterior distribution. MCMC chains are initialised for each of the distributions with occasional swaps between the distributions. The sample collected at the lowest temperature is a sample from the posterior distribution. Parallel tempering allows larger (between modes) moves made at the higher temperature distributions to be passed down to the lower temperature distributions. Let $\boldsymbol{\delta} = (\boldsymbol{\theta}, \boldsymbol{\Omega}, \rho, \mathbf{d}^*, \boldsymbol{\Sigma}, \boldsymbol{\psi}, \mathbf{y}_C)$ be the collection of unknown parameters and censored observations. For $r = 1, \ldots, R$, define the rth distribution to have density proportional to

$$U_r(\boldsymbol{\delta}) = \exp\left(\frac{\log \pi(\mathbf{y}_S|\boldsymbol{\delta}) + \log \pi(\boldsymbol{\delta})}{\tau_r}\right),$$

where $1 = \tau_1 < \cdots < \tau_R$ are a series of "temperatures".

A parallel tempering iteration involves applying either a a) sampling step (with probability ω); or b) swap step (with probability $1-\omega$). In the sampling step, the current values $\boldsymbol{\delta}^c_{(r)}$ in the rth chain are updated to $\boldsymbol{\delta}^{c+1}_{(r)}$, for $r = 1, \ldots, R$ using the Gibbs sampling algorithm. In the swap step, two neighbouring chains (r and $r+1$) are chosen at random. With probability

$$\min\left[1, \frac{U_{r+1}(\boldsymbol{\delta}^c_{(r)})U_r(\boldsymbol{\delta}^c_{(r+1)})}{U_{r+1}(\boldsymbol{\delta}^c_{(r+1)})U_r(\boldsymbol{\delta}^c_{(r)})}\right]$$

set $\boldsymbol{\delta}^{c+1}_{(r)} = \boldsymbol{\delta}^c_{(r+1)}$ and $\boldsymbol{\delta}^{c+1}_{(r+1)} = \boldsymbol{\delta}^c_{(r)}$. To complete the sampling step (a) for each distribution at temperature t_r, the full conditional and RMLMH components need to be adjusted (see Appendices A and B for details).

3.2 Amended Gibbs sampling and parallel tempering

One iteration of the RMLMH step in the Gibbs sampling algorithm outlined in Section 3.1 requires one evaluation of $\mathbf{m}(\boldsymbol{\theta})$ to calculate the acceptance probability and one evaluation each of $\partial\mathbf{m}(\boldsymbol{\theta})/\partial\boldsymbol{\theta}$ and $\partial^2\mathbf{m}(\boldsymbol{\theta})/\partial\boldsymbol{\theta}\partial\theta_p$ to form the proposal distribution. To evaluate $\mathbf{m}(\boldsymbol{\theta})$ we are required to evaluate $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x}_i,t_{ij})$, for $i=1,\ldots,I$ and $j=1,\ldots,n_i$. The derivative terms require evaluation of $\partial\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x}_i,t_{ij})/\partial\boldsymbol{\theta}$ and $\partial^2\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x}_i,t_{ij})/\partial\boldsymbol{\theta}\partial\theta_k$, for $i=1,\ldots,I$, $j=1,\ldots,n_i$ and $k=1,\ldots,p$. We can augment the system of ODEs with additional equations to numerically solve for the first- and second-order sensitivities (Girolami and Calderhead, 2011). However this will substantially increase the computational burden of the numerical solution. Note that to generate a sample of size B using parallel tempering with R chains will require BR evaluations of the three functions for each of $\mathbf{x}_i$ and t_{ij}, for $i=1,\ldots,I$, $j=1,\ldots,n_i$ and $k=1,\ldots,p$.

Once we have generated a sample of size B from the posterior distribution, to estimate the probability $\mathrm{P}(\mathbf{y}_0\in\mathcal{Y}|\mathbf{y}_S)$, we require B further evaluations of $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)$, for each $\mathbf{x}\in\mathcal{X}$ and $t\in\mathcal{T}$ at which we wish to evaluate this probability.

Making such a large number of evaluations of the numerical solution is clearly computationally infeasible. Instead, prior to each of the two phases described in Section 2.5, we construct a multivariate Gaussian process (MGP) emulator and employ this emulator to reduce the computational burden. Statistical emulators are commonly-used to predict the output of computationally expensive functions. They are based on a "small" number, N, of function evaluations and provide a predictive distribution with which to approximate the function of interest. The predictive mean (or mode, median, etc) can be used as a point prediction and the distribution allows uncertainty in the prediction to be quantified. Let $\mathrm{Q}(\boldsymbol{\theta},\mathbf{x},t)$ denote the predictive distribution with predictive mean denoted by $\hat{\boldsymbol{\mu}}(\boldsymbol{\theta};\mathbf{x},t)$.

In the Sampling Phase, we apply the Gibbs sampling and parallel tempering algorithm with the following amendments:

1. All evaluations of $\partial\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)/\partial\boldsymbol{\theta}$ and $\partial^2\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)/\partial\boldsymbol{\theta}\partial\theta_k$ are replaced by $\partial\hat{\boldsymbol{\mu}}(\boldsymbol{\theta};\mathbf{x},t)/\partial\boldsymbol{\theta}$ and $\partial^2\hat{\boldsymbol{\mu}}(\boldsymbol{\theta};\mathbf{x},t)/\partial\boldsymbol{\theta}\partial\theta_k$, respectively, for $k=1,\ldots,p$.

2. Construct $R+1$ parallel chains with temperatures $1=\tau_0=\tau_1<\tau_2<\cdots<\tau_R$. In chains $r\geq 1$, all evaluations of $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)$ in the evaluation of $U_r(\boldsymbol{\delta})$ in the acceptance probability are replaced by evaluation of $\hat{\boldsymbol{\mu}}(\boldsymbol{\theta};\mathbf{x},t)$. For evaluation of $U_0(\boldsymbol{\delta})$, true evaluation of $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)$ is retained.

See Fielding et al. (2011) for a similar parallel tempering scheme.

Therefore the sample generated in chain $r=0$ is from the true posterior distribution of $\boldsymbol{\delta}$ but the sampling scheme only requires B evaluations of $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x}_i,t_{ij})$ for $i=1,\ldots,I$ and $j=1,\ldots,n_i$.

In the Prediction Phase, to estimate the probability of satisfying the constraints, all evaluations of $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)$ are replaced by a value generated from $\mathrm{Q}(\boldsymbol{\theta},\mathbf{x},t)$. This means the estimate of the probability takes account of the uncertainty in the prediction made by the emulator.

We describe how each of the emulators, prior to the Sampling and Prediction Phases, are formed in Section 3.3.

3.3 Multivariate Gaussian process emulator

In this paper we use the MGP (Conti and O'Hagan, 2010; Overstall and Woods, 2016). An MGP is formed by evaluating $\boldsymbol{\mu}(\boldsymbol{\theta};\mathbf{x},t)$ at a carefully selected set (or meta-design) $\zeta=\{(\boldsymbol{\theta}_1^*,\mathbf{x}_1^*,t_1^*),\ldots,(\boldsymbol{\theta}_N^*,\mathbf{x}_N^*,t_N^*)\}$. The superscript * notation has been introduced to distinguish the meta-design from the design of the physical experiment. Let $\mathbf{z}_i=\boldsymbol{\mu}(\boldsymbol{\theta}_i^*;\mathbf{x}_i^*,t_i^*)$, $\mathbf{w}_i=(\boldsymbol{\theta}_i^*;\mathbf{x}_i^*,t_i^*)$ be the $U\times 1$ vector giving all inputs, where $U=11$, and $\mathbf{Z}$ be the $N\times K$ matrix with ith row given by $\mathbf{z}_i^T$. The central assumption of the MGP is

$$\mathbf{Z}\sim\mathrm{MN}\left(\mathbf{1}_N\boldsymbol{\beta}^T,\boldsymbol{\Phi},\mathbf{P}\right),$$

where $\boldsymbol{\beta}$ is the $K \times 1$ vector giving each of the common column means of $\mathbf{Z}$, $\boldsymbol{\Phi}$ is a $K \times K$ unstructured column covariance matrix and $\mathbf{P}$ is an $N \times N$ row correlation matrix. The ijth element of $\mathbf{P}$ is given by

$$P_{ij} = \exp\left(-\sum_{u=1}^{U} \alpha_u \left(w_{iu} - w_{ju}\right)^2\right),$$

where w_{iu} is the uth element of $\mathbf{w}_i$, for $u = 1, \ldots, U$ and $\alpha_u > 0$, for $u = 1, \ldots, U$, are correlation parameters. Suppose we wish to predict $\mathbf{z}_0 = \boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$. Let

$$\hat{\boldsymbol{\beta}} = \frac{\mathbf{1}_N^T \mathbf{P}^{-1} \mathbf{Z}}{\mathbf{1}_N^T \mathbf{P}^{-1} \mathbf{1}_N}, \tag{14}$$

$$\hat{\boldsymbol{\Psi}} = \mathbf{Z}^T \mathbf{P}^{-1} \left(\mathbf{I}_N - \frac{\mathbf{1}_N \mathbf{1}_N^T \mathbf{P}^{-1}}{\mathbf{1}_N^T \mathbf{P}^{-1} \mathbf{1}_N}\right) \mathbf{Z}, \tag{15}$$

and $\mathbf{w} = (\boldsymbol{\theta}, \mathbf{x}, t)$. It can be shown that

$$\mathrm{Q}(\boldsymbol{\theta}, \mathbf{x}, t) = \mathrm{N}\left(\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t), \hat{\boldsymbol{\Pi}}(\boldsymbol{\theta}; \mathbf{x}, t)\right), \tag{16}$$

where

$$\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t) = \hat{\boldsymbol{\beta}} + \mathbf{p}^T \mathbf{P}^{-1} \left(\mathbf{Z} - \mathbf{1}_N \hat{\boldsymbol{\beta}}^T\right), \tag{17}$$

$$\hat{\boldsymbol{\Pi}}(\boldsymbol{\theta}; \mathbf{x}, t) = \left(1 - \mathbf{p}^T \mathbf{P}^{-1} \mathbf{p}\right) \hat{\boldsymbol{\Psi}}, \tag{18}$$

with $\mathbf{p}$ having ith element

$$p_i = \exp\left(-\sum_{u=1}^{U} \alpha_u \left(w_{iu} - w_u\right)^2\right),$$

and w_u being the uth element of $\mathbf{w}$. The predictive distribution given by (16) is conditional on the correlation parameters, $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_U)$. We replace these values by their marginal posterior mode (e.g. Overstall and Woods, 2016) although a fully Bayesian approach could be taken.

The two emulators formed prior to the Sampling and Prediction Phases are essentially defined by the choice of the meta-design ζ. The numerical solution, $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$, has three types of input: $\boldsymbol{\theta}$, $\mathbf{x}$ and t. In the choice of ζ, we treat these three types of input separately for two reasons. Firstly, when we choose ζ for the emulator for the Sampling Phase, it makes sense for the values of $\mathbf{x}$ and t to be those of the design of the physical experiment, i.e. $\mathbf{x}_1, \ldots, \mathbf{x}_I$ and $t_1, \ldots, t_m$. This means that the uncertainty in the prediction will only be over Θ, the parameter space of $\boldsymbol{\theta}$. Secondly, by treating the input t separately, we can exploit the efficiency of computing the numerical solution for the same $\boldsymbol{\theta}$ and $\mathbf{x}$ but for different t, i.e. to compute $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t_1)$ and $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t_2)$, for $t_1 \neq t_2$, requires a negligible increase in computational effort over just computing $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t_1)$.

Therefore, we let the meta-design be

$$\zeta = \zeta_1 \times \zeta_2 \times \zeta_3,$$

where, in this case, $\times$ denotes the Cartesian product, $\zeta_1 = \left\{\boldsymbol{\theta}_1^*, \ldots, \boldsymbol{\theta}_{N_1}^*\right\}$, $\zeta_2 = \left\{\mathbf{x}_1^*, \ldots, \mathbf{x}_{N_2}^*\right\}$, $\zeta_3 = \left\{t_1^*, \ldots, t_{N_3}^*\right\}$, and $N = \prod_{i=1}^{M} N_i$ where $M = 3$.

A further advantage of this structure is that the row correlation matrix, $\mathbf{P}$, can be written as

$$\mathbf{P} = \bigotimes_{i=1}^{M} \mathbf{P}_i = \mathbf{P}_1 \otimes \mathbf{P}_2 \otimes \mathbf{P}_3,$$

where the ijth elements of $\mathbf{P}_1$, $\mathbf{P}_2$ and $\mathbf{P}_3$ are

$$\begin{aligned}
P_{1,ij} &= \exp\left(-\textstyle\sum_{u=1}^{p} \alpha_u \left(\theta_{iu}^* - \theta_{ju}^*\right)^2\right) && \text{for } i, j = 1, \ldots, N_1, \\
P_{2,ij} &= \exp\left(-\textstyle\sum_{u=1}^{F} \alpha_{p+u} \left(x_{iu}^* - x_{ju}^*\right)^2\right) && \text{for } i, j = 1, \ldots, N_2, \\
P_{3,ij} &= \exp\left(-\alpha_U \left(t_i^* - t_j^*\right)^2\right) && \text{for } i, j = 1, \ldots, N_3,
\end{aligned}$$

respectively. This means that

$$\mathbf{P}^{-1} = \bigotimes_{i=1}^{M} \mathbf{P}_i^{-1},$$

which significantly reduces the computational burden and improves the numerical stability of inverting $\mathbf{P}^{-1}$. Similarly, (14), (15), (17) and (18) simplify to

$$\begin{aligned}
\hat{\boldsymbol{\beta}} &= \frac{\left(\bigotimes_{i=1}^{M} \mathbf{1}_{N_i}^T \mathbf{P}_i^{-1}\right) \mathbf{Z}}{\prod_{i=1}^{M} \mathbf{1}_{N_i}^T \mathbf{P}_i^{-1} \mathbf{1}_{N_i}}, \\
\hat{\boldsymbol{\Psi}} &= \mathbf{Z}^T \left[\bigotimes_{i=1}^{M} \mathbf{P}_i^{-1} - \bigotimes_{i=1}^{M} \mathbf{P}_i^{-1} \mathbf{1}_{N_i} \mathbf{1}_{N_i}^T \mathbf{P}_i^{-1}\right] \mathbf{Z}, \\
\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t) &= \hat{\boldsymbol{\beta}} + \left(\bigotimes_{i=1}^{M} \mathbf{p}_i^T \mathbf{P}_i^{-1}\right) \left(\mathbf{Z} - \mathbf{1}_N \hat{\boldsymbol{\beta}}^T\right), \\
\hat{\boldsymbol{\Pi}}(\boldsymbol{\theta}; \mathbf{x}, t) &= \left(1 - \prod_{i=1}^{M} \mathbf{p}_i^T \mathbf{P}_i^{-1} \mathbf{p}_i\right) \hat{\boldsymbol{\Psi}},
\end{aligned}$$

respectively, where $\mathbf{p}_i$ (for $i = 1, \ldots, M$) has jth element given by

$$\begin{aligned}
p_{1j} &= \exp\left(-\textstyle\sum_{u=1}^{p} \alpha_u \left(\theta_{ju}^* - \theta_u\right)^2\right), && \text{for } j = 1, \ldots, N_1, \\
p_{2j} &= \exp\left(-\textstyle\sum_{u=1}^{p} \alpha_{p+u} \left(x_{ju}^* - x_u\right)^2\right), && \text{for } j = 1, \ldots, N_2, \\
p_{3j} &= \exp\left(-\alpha_U \left(t_j^* - t\right)^2\right), && \text{for } j = 1, \ldots, N_3.
\end{aligned}$$

3.3.1 Meta-design for the Sampling Phase

Fix $\zeta_2 = \{\mathbf{x}_1, \ldots, \mathbf{x}_5\}$ and $\zeta_3 = \{t_1, \ldots, t_m\}$, i.e. $N_2 = 5$ and $N_3 = m$. We use the following exploratory algorithm to adaptively choose ζ_1. This uses a hybrid of the methodologies proposed by Rasmussen (2003), Fielding et al. (2011) and Overstall and Woods (2013). Rasmussen (2003) first proposed using MCMC methods to adaptively improve an emulator to an unnormalised posterior density based on a computationally expensive function in light of observed data. Fielding et al. (2011) extended this methodology by using parallel tempering. Overstall and Woods (2013) proposed to emulate the computationally expensive function, as opposed to the posterior density, to make model criticism more efficient.

The idea of exploratory algorithm is to iteratively improve the accuracy of the emulator in the region of Θ corresponding to high posterior density.

1. Generate the set $\zeta_1^0 = \left\{\boldsymbol{\theta}_1^*, \ldots, \boldsymbol{\theta}_{N_1^0}^*\right\}$ as a sample of size N_1^0 from the prior distribution of $\boldsymbol{\theta}$.
2. Let $\zeta^0 = \zeta_1^0 \times \zeta_2 \times \zeta_3$ and fit the MGP to the resulting evaluations from $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$.
3. Let $\zeta = \zeta^0$ and repeat the following steps until $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$ as been evaluated a total of N times.
   (a) Perform L iterations of the Gibbs sampling and parallel tempering scheme (for chains $r = 1, \ldots, R$) where evaluation of $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$ is replaced by evaluation of $\hat{\boldsymbol{\mu}}(\boldsymbol{\theta}; \mathbf{x}, t)$ in all instances.
   (b) Evaluate $\tilde{\mathbf{z}}_{rij} = \boldsymbol{\mu}(\boldsymbol{\theta}_{(r)}; \mathbf{x}_i, t_j)$, for $r = 1, \ldots, R$, $i = 1, \ldots, I$ and $j = 1, \ldots, n_i$, where $\boldsymbol{\theta}_{(r)}$ is the current value of $\boldsymbol{\theta}$ in the rth chain. Augment the matrix $\mathbf{Z}$ by $\tilde{\mathbf{z}}_{rij}$ for $r = 1, \ldots, R$, $i = 1, \ldots, I$ and $j = 1, \ldots, n_i$. Refit the MGP.

3.3.2 Meta-design for the Prediction Phase

Let ζ_1 be a space-filling sub-sample (e.g. Johnson et al., 1990) of size N_1 from the sample generated from the posterior distribution of $\boldsymbol{\theta}$ in the Sampling Phase. Furthermore, let ζ_2 and ζ_3 be Latin hypercube designs in $\mathcal{X}$ and $\mathcal{T}$ of sizes N_2 and N_3, respectively.

4 Results

4.1 Sampling Phase

To create the meta-design for the Sampling Phase, we set $N_1^0 = 50$, $N_1 = 100$ and set up $R = 5$ parallel chains. This means we perform ten iterations of the exploratory algorithm in Section 3.3.1 with $L = 50$. We adapt the temperatures of the parallel chains using the methodology of Miasojedow et al. (2013) as implemented by Overstall and Woods (2013). Figure 1 shows the values of $\boldsymbol{\theta}^*$ in the meta-design ζ_1. The grey points are the N_1^0 initial values generated from the prior distribution and the black points are the values selected in step 3 of the exploratory algorithm. In the latter case, it is clear that these points are being sampled from a concentrated region of Θ.

We then obtain a sample of size $B = 50000$ from the posterior distribution of $\boldsymbol{\delta}$ using the amended Gibbs sampling and parallel tempering algorithm presented in Section 3.2. Figure 2 shows trace plots of the posterior sample (left hand panels) and the estimated posterior density (right hand panels), for each element of $\boldsymbol{\theta}$. The right hand panels also show the density of the prior distribution (in grey).

Diagnostics for assessing the adequacy of the fitted model (see, for example, Gelman et al., 2014, Chapter 6) indicated that there were no reasons to believe that the model was inadequate.

4.2 Prediction Phase

We now use the sample generated from the posterior distribution of $\boldsymbol{\delta}$, and, in particular, $\boldsymbol{\theta}$, to construct an emulator for the numerical solution to the ODEs. We let $N_1 = 20$, $N_2 = 50$ and $N_3 = 50$. It could be noted that this value of N_1 is quite small compared to the rule of thumb of Loeppky et al. (2009) that the sample size should be approximately ten times the number of input dimensions (in this case $p = 5$). However, we found that, due to the concentration of the posterior, this value was sufficiently large to produce an adequate emulator and avoid numerical instability in the inversion of $\mathbf{P}_1$. To assess the adequacy of the emulator, we created a test design of the same size and implemented the diagnostic methods of Overstall and Woods (2016). In particular, the mean squared errors between the numerical solution and the predictive mean were 1.5×10^{-3}, 0.015 and 0.076 for each of the $K = 3$ dimensions of $\boldsymbol{\mu}(\boldsymbol{\theta}; \mathbf{x}, t)$. The overall coverage of the 95% predictive intervals was 94.2%.

Table 3: Optimum values of the controllable factors, $\mathbf{x}$ and t, as found by ACE.

Name	Symbol		Value	Unit
Initial amount of A	A_0	x_1	39.09	mols
Initial amount of D	D_0	x_2	91.59	mols
Initial amount of E	E_0	x_3	26.47	mols
Temperature	λ	x_4	40.00	°C
Volume	V	x_5	31.28	l
Time	t		180.60	s

We now have a stochastic approximation to the probability, $\mathrm{P}(\mathbf{y}_0 \in \mathcal{Y} | \mathbf{y}_S)$, of satisfying the constraints for given $\mathbf{x}$ and t. We need to maximise this probability over the space $\mathcal{S} = \mathcal{X} \times \mathcal{T}$. To do this we use the approximate coordinate exchange (ACE; Overstall and Woods 2015) algorithm. This can be used to maximise

Figure 1: Values of $\boldsymbol{\theta}$ in the meta-design ζ_1 for the Sampling Phase as found by the exploratory algorithm in Section 3.3.1.

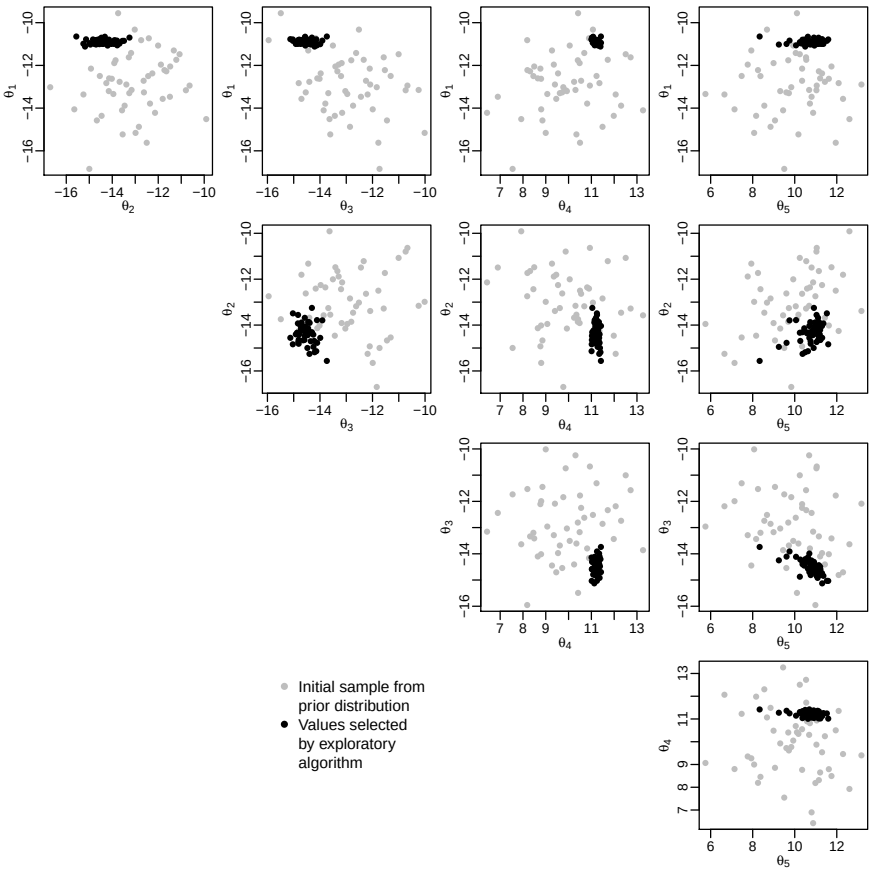

Figure 2: Trace plots of the posterior sample (left hand panels) and the estimated posterior and prior densities (right hand panels) for each element of $\boldsymbol{\theta}$

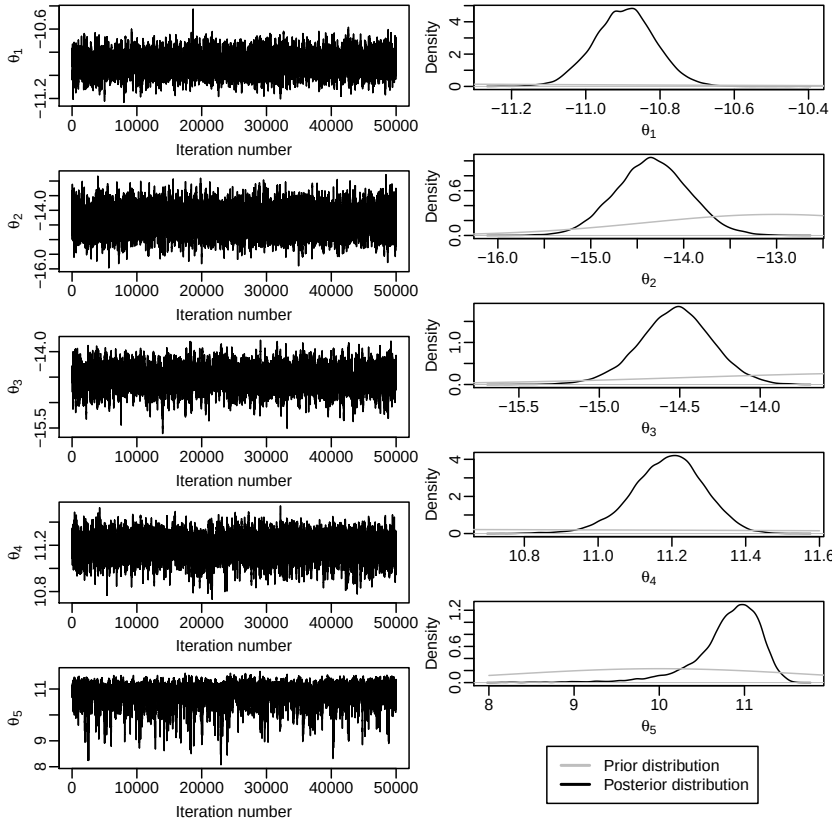

Figure 3: The probability of satisfying each of the three marginal constraints and the joint constraints for time (in s; x-axis) against initial amount of A (in mols; y-axis) for five different values of temperature. The colours indicate the magnitude of the probability, with black being 0 and white being 1. The white circle in the 5th column and last row indicates the approximately optimal set of controllable factors.

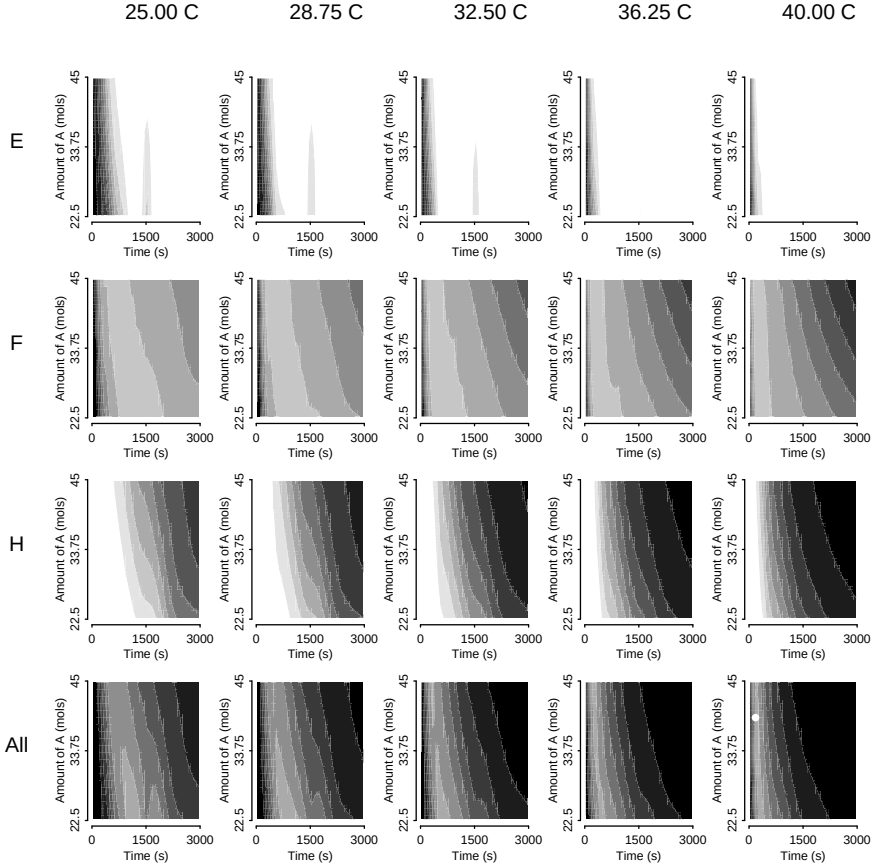

stochastic approximations to objective functions by using the cyclic ascent algorithm in conjunction with a univariate Gaussian process emulator approximation to the objective function. We use 100 random starts of ACE of twenty iterations each, where the initial value of $(\mathbf{x}, t)$ is uniformly generated in $\mathcal{S}$. The optimum values of $\mathbf{x}$ and t, as found by ACE, are shown in Table 3. The approximate (subject to Monte Carlo error) maximum value of $\mathrm{P}(\mathbf{y}_0 \in \mathcal{Y}|\mathbf{y}_S)$ attained was 0.70. To further explore how $\mathrm{P}(\mathbf{y}_0 \in \mathcal{Y}|\mathbf{y}_S)$ depends on $\mathbf{x}$ and t, we fix the initial amounts of D, E and Volume at 91.59 mols, 26.47 mols and 31.28 l, respectively, i.e. the optimum values as found by ACE. We then vary the initial amounts of A, the temperature and time over the ranges identified in Table 1. The last row of Figure 3 shows $\mathrm{P}(\mathbf{y}_0 \in \mathcal{Y}|\mathbf{y}_S)$ for time against the initial amount of A, with different columns for different temperatures. The first three rows show corresponding plots for the marginal probability of satisfying each of the three constraints (3) on E, F and H, respectively. The trade-off between the objectives given by the three constraints can be clearly seen. To satisfy the constraint $E(t) < 3$ mols, t has to be large. This is intuitively obvious since the initial amount of E is 26.47 mols so the process will need to progress for some time before the amount of E has decreased enough to satisfy the constraint. The opposite is true for H (initial amount of 0 mols), where the optimum time to satisfy $H(t) < 3$ mols is for t to be small. This means that there is a very narrow window of values of t that will satisfy all three constraints, as shown in the last row of Figure 3. The approximately optimal set of controllable factors is identified in Figure 3 by the white circle in the 5th column and last row.

5 Discussion

This paper has considered specifying the controllable factors of a chemical synthesis process used to produce a pharmaceutical product, to ensure that the final amounts of certain substance respect certain constraints, specifically that a) the amount of the substance constituting the pharmaceutical product is above some level; and b) the amounts of two unwanted by-product substances are below some levels. The relationship between the controllable factors and the amounts of substances of interest is hypothesised to be governed by the analytically intractable solution to a system of ODEs. Responses from a physical experiment are used to refine the knowledge on this relationship and the probability of satisfying the constraints is maximised. Statistical emulators for the numerical solution are constructed to accelerate the model-fitting process and the estimation of the probability of satisfying the constraints.

A Full-conditional distributions for Gibbs sampling

A.1 Full conditional distribution for model discrepancy

For the rth chain with temperature τ_r, the full-conditional distribution of $\mathbf{d}^*$ is

$$\mathbf{d}^*|\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}, \rho, \psi \sim \mathrm{N}\left(\boldsymbol{\mu}_{\mathbf{d}^*}, \tau_r \mathbf{V}_{\mathbf{d}^*}\right),$$

where

$$\begin{aligned} \mathbf{V}_{\mathbf{d}^*} &= \left(\mathbf{H}^T \left(\boldsymbol{\Omega}^{-1} \otimes \mathbf{T}^{-1}\right) \mathbf{H} + \left(\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}\right)\right)^{-1}, \\ \boldsymbol{\mu}_{\mathbf{d}^*} &= \mathbf{V}_{\mathbf{d}^*} \left(\mathbf{H}^T \left(\boldsymbol{\Omega}^{-1} \otimes \mathbf{T}^{-1}\right) \mathbf{y} + \left(\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}\right) \mathbf{m}(\boldsymbol{\theta})\right). \end{aligned}$$

A.2 Full conditional distributions for covariance matrices

For the rth chain with temperature τ_r, the full-conditional distributions of $\boldsymbol{\Omega}$ and $\boldsymbol{\Sigma}$ are given by

$$\boldsymbol{\Omega}|\mathbf{Y}, \mathbf{D}^*, \rho \sim \mathrm{IW}\left(\frac{\nu + n + K + 1}{\tau_r} - K - 1, \frac{1}{\tau_r}\left(\mathbf{I}_K + \left(\mathbf{Y} - \mathbf{G}\mathbf{D}^*\right)^T \mathbf{T}^{-1} \left(\mathbf{Y} - \mathbf{G}\mathbf{D}^*\right)\right)\right),$$

and

$$\boldsymbol{\Sigma}|\mathbf{D}^*, \boldsymbol{\theta}, \psi \sim \mathrm{IW}\left(\frac{\nu + m + K + 1}{\tau_r} - K - 1, \frac{1}{\tau_r}\left(\mathbf{I}_k + \left(\mathbf{D}^* - \mathbf{M}(\boldsymbol{\theta})\right)^T \mathbf{S}^{-1} \left(\mathbf{D}^* - \mathbf{M}(\boldsymbol{\theta})\right)\right)\right),$$

respectively, where IW denotes the inverse-Wishart distribution.

A.3 Full conditional distribution for censored observations

Let $\mathbf{A}$ be the $n \times n$ permutation matrix which re-orders the elements of $\mathbf{y}$ such that we can write

$$\mathbf{A}\mathbf{y} = \begin{pmatrix} \mathbf{y}_S \\ \mathbf{y}_C \end{pmatrix} \sim \mathrm{N}\left(\mathbf{A}\mathbf{H}\mathbf{d}^*, \mathbf{A}\left(\boldsymbol{\Omega} \otimes \mathbf{T}\right)\mathbf{A}^T\right).$$

Define $\mathbf{b} = \mathbf{A}\mathbf{H}\mathbf{d}^*$ and $\mathbf{L} = \mathbf{A}\left(\boldsymbol{\Omega} \otimes \mathbf{T}\right)\mathbf{A}^T$. Let $\mathbf{b} = \left(\mathbf{b}_S, \mathbf{b}_c\right)^T$ where $\mathbf{b}_S$ and $\mathbf{b})C$ are the $n_S \times 1$ and $n_c \times 1$ sub-vectors corresponding to $\mathbf{y}_S$ and $\mathbf{y}_C$, respectively. Similarly, partition $\mathbf{L}$ as

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}_{SS} & \mathbf{L}_{SC} \\ \mathbf{L}_{CS} & \mathbf{L}_{CC} \end{pmatrix}.$$

The full conditional distribution of $\mathbf{y}_c$ is

$$\mathbf{y}_C \sim \mathrm{N}\left(\mathbf{b}_C + \mathbf{L}_{CS}\mathbf{L}_{SS}^{-1}\left(\mathbf{y}_S - \mathbf{b}_S\right), \tau_r\left(\mathbf{L}_{CC} - \mathbf{L}_{CS}\mathbf{L}_{SS}^{-1}\mathbf{L}_{SC}\right)\right)$$

truncated to the hypercube $[-\infty, \log\chi]^{n_c}$.

B Riemann manifold Langevin Metropolis-Hastings (RMLMH) Algorithm

Consider generating a sample from the distribution of $\boldsymbol{\delta} \in \mathbb{R}^p$, conditional on $\mathbf{y}$, having density $\pi(\boldsymbol{\delta}|\mathbf{y}) \propto (\pi(\mathbf{y}|\boldsymbol{\delta})\pi(\boldsymbol{\delta}))^{1/\tau_r}$, for temperature τ_r, using the Metropolis-Hastings algorithm. Let $h(\boldsymbol{\delta}) = \log \pi(\mathbf{y}|\boldsymbol{\delta}) + \log \pi(\boldsymbol{\delta})$ and define

$$\begin{aligned} \bigtriangledown h(\boldsymbol{\delta}) &= \frac{\partial \log \pi(\mathbf{y}|\boldsymbol{\delta})}{\partial \boldsymbol{\delta}} + \frac{\partial \log \pi(\boldsymbol{\delta})}{\partial \boldsymbol{\delta}}, \\ \mathbf{G}(\boldsymbol{\delta}) &= -\mathrm{E}_{\mathbf{y}|\boldsymbol{\delta}}\left[\frac{\partial \log \pi(\mathbf{y}|\boldsymbol{\delta})}{\partial \boldsymbol{\delta}} \frac{\partial \log \pi(\mathbf{y}|\boldsymbol{\delta})}{\partial \boldsymbol{\delta}^T}\right] - \frac{\partial^2 \log \pi(\boldsymbol{\delta})}{\partial \boldsymbol{\delta} \partial \boldsymbol{\delta}^T}. \end{aligned}$$

If $\boldsymbol{\delta}^c$ is the current value of $\boldsymbol{\delta}$, then under the RMLMH algorithm, the proposal distribution is $\mathrm{N}\left(\boldsymbol{\mu}_{\boldsymbol{\delta}}, \tau_r \mathbf{V}_{\boldsymbol{\delta}}\right)$, where

$$\begin{aligned} \boldsymbol{\mu}_{\boldsymbol{\delta}} &= \boldsymbol{\delta}^c + \frac{\epsilon^2}{2}\mathbf{G}(\boldsymbol{\delta}^c)^{-1} \bigtriangledown h(\boldsymbol{\delta}^c) - \epsilon^2 \mathbf{n}^{(1)} + \frac{\epsilon^2}{2}\mathbf{n}^{(2)}, \\ \mathbf{V}_{\boldsymbol{\delta}} &= \epsilon^2 \mathbf{G}(\boldsymbol{\delta}^c)^{-1}, \end{aligned}$$

with $\mathbf{n}^{(1)}$ and $\mathbf{n}^{(2)}$ having ith elements

$$\begin{aligned} n_i^{(1)} &= \sum_{j=1}^{p} \left\{ \mathbf{G}(\boldsymbol{\delta}^c)^{-1} \left. \frac{\partial \mathbf{G}(\boldsymbol{\delta})}{\partial \delta_j} \right|_{\boldsymbol{\delta}=\boldsymbol{\delta}^c} \mathbf{G}(\boldsymbol{\delta}^c)^{-1} \right\}_{ij}, \\ n_i^{(2)} &= \sum_{j=1}^{p} \left\{ \mathbf{G}(\boldsymbol{\delta}^c)^{-1} \right\}_{ij} \mathrm{tr}\left\{ \mathbf{G}(\boldsymbol{\delta}^c)^{-1} \left. \frac{\partial \mathbf{G}(\boldsymbol{\delta})}{\partial \delta_j} \right|_{\boldsymbol{\delta}=\boldsymbol{\delta}^c} \right\}, \end{aligned}$$

respectively.

C Components required for Riemann manifold MH Algorithm

C.1 Physical parameters

The log density of the full conditional distribution of $\boldsymbol{\theta}$ is given by

$$h(\boldsymbol{\theta}) \propto -\frac{1}{2}\left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right)^T \left(\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}\right)\left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right) - \frac{1}{2}\left(\boldsymbol{\theta} - \boldsymbol{\mu}\right)^T \boldsymbol{\Delta}^{-1}\left(\boldsymbol{\theta} - \boldsymbol{\mu}\right).$$

The derivative with respect to $\boldsymbol{\theta}$ is

$$\bigtriangledown h(\boldsymbol{\theta}) = \left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right)^T \left(\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}\right) \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} - \boldsymbol{\Delta}^{-1}\left(\boldsymbol{\theta} - \boldsymbol{\mu}\right).$$

The matrix tensor for $\boldsymbol{\theta}$ is

$$\mathbf{G}(\boldsymbol{\theta}) = \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T}\left(\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}\right) \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} + \boldsymbol{\Delta}^{-1},$$

with derivatives

$$\frac{\partial \mathbf{G}(\boldsymbol{\theta})}{\partial \theta_k} = 2\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T}\left(\boldsymbol{\Sigma}^{-1} \otimes \mathbf{S}^{-1}\right) \frac{\partial^2 \mathbf{m}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \theta_k},$$

for $k = 1, \ldots, p$.

C.2 Correlation parameter for time dependency

Consider a log transformation of ρ, i.e. $a = \log \rho$. The log density of the full conditional distribution of a is given by

$$\begin{aligned} h(a) \quad \propto \quad & -\frac{K}{2}\sum_{i=1}^{I'} \log|\mathbf{T}_i| - \frac{1}{2}\left(\mathbf{y}-\mathbf{H}\mathbf{d}^*\right)^T \left(\mathbf{\Omega}^{-1} \otimes \text{diag}_{i=1,\ldots,I'}\left\{\mathbf{T}_i^{-1}\right\}\right)\left(\mathbf{y}-\mathbf{H}\mathbf{d}^*\right) \\ & + a - \exp(a). \end{aligned}$$

The derivative with respect to a is

$$\begin{aligned} \frac{\partial h(a)}{\partial a} \quad = \quad & -\frac{K}{2}\sum_{i=1}^{I'} \text{tr}\left\{\mathbf{T}_i^{-1}\mathbf{T}_{ia}\right\} \\ & + \frac{1}{2}\left(\mathbf{y}-\mathbf{H}\mathbf{d}^*\right)^T \left(\mathbf{\Omega}^{-1} \otimes \text{diag}_{i=1,\ldots,I'}\left\{\mathbf{T}_i^{-1}\mathbf{T}_{ia}\mathbf{T}_i^{-1}\right\}\right)\left(\mathbf{y}-\mathbf{H}\mathbf{d}^*\right) \\ & + 1 - \exp(a), \end{aligned}$$

where $\mathbf{T}_{ia} = \partial \mathbf{T}_i / \partial a$ with jlth element $-\exp(a)(t_{ij}-t_{il})^2 T_{i,jl}$. The tensor matrix for a is

$$\mathbf{G}(a) = \frac{K}{2}\text{tr}\left\{\mathbf{T}_i^{-1}\mathbf{T}_{ia}\mathbf{T}_i^{-1}\mathbf{T}_{ia}\right\} - \exp(a),$$

with derivative

$$\frac{\partial \mathbf{G}(a)}{\partial a} = K\text{tr}\left\{\mathbf{T}_i^{-1}\mathbf{T}_{ia}\mathbf{T}_i^{-1}\left(\mathbf{T}_{iaa} - \mathbf{T}_{ia}\mathbf{T}_i^{-1}\mathbf{T}_{ia}\right)\right\} - \exp(a),$$

where $\mathbf{T}_{iaa} = \partial^2 \mathbf{T}_i / \partial a^2$ with jlth element $\exp(a)(t_{ij}-t_{il})^2\left(\exp(a)(t_{ij}-t_{il})^2 - 1\right)T_{i,jl}$.

C.3 Correlation parameters for model discrepancy

Consider a log transformation of each element of $\boldsymbol{\psi}$, i.e. $b_i = \log \psi_i$ for $i = 1, 2$. The log density of the full conditional distribution of $\mathbf{b} = (b_1, b_2)$ is given by

$$h(\mathbf{b}) \propto -\frac{K}{2}\log|\mathbf{S}| - \frac{1}{2}\left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right)^T \left(\mathbf{\Sigma} \otimes \mathbf{S}^{-1}\right)\left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right) + \sum_{i=1}^{2} b_i - \exp(b_i).$$

The derivative with respect to b_i is

$$\frac{\partial h(\mathbf{b})}{\partial b_i} = -\frac{K}{2}\text{tr}\left\{\mathbf{S}^{-1}\mathbf{S}_i\right\} + \frac{1}{2}\left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right)^T \left(\mathbf{\Sigma} \otimes \mathbf{S}^{-1}\mathbf{S}_i\mathbf{S}^{-1}\right)\left(\mathbf{d}^* - \mathbf{m}(\boldsymbol{\theta})\right) + 1 - \exp(b_i),$$

where $\mathbf{S}_i = \partial \mathbf{S} / \partial b_i$ with jlth element given by

$$\begin{cases} -\exp(b_1)\left[\sum_{f=1}^{F}\left(d_{jf}-d_{lf}\right)^2\right] S_{jl}, & \text{if } i = 1, \\ -\exp(b_2)(t_j - t_l)^2 S_{jl}, & \text{if } i = 2. \end{cases}$$

The tensor matrix, $\mathbf{G}(\mathbf{b})$ has ijth element

$$G(\mathbf{b})_{ij} = \frac{K}{2}\text{tr}\left\{\mathbf{S}^{-1}\mathbf{S}_i\mathbf{S}^{-1}\mathbf{S}_j\right\} - I(i=j)\exp(b_i).$$

The derivatives of $\mathbf{G}(\mathbf{b})$ are

$$\begin{aligned} \frac{\partial G(\mathbf{b})_{ij}}{\partial b_k} \quad = \quad & \frac{K}{2}\text{tr}\left\{\mathbf{S}^{-1}\mathbf{S}_i\mathbf{S}^{-1}\mathbf{S}_{jk} - \mathbf{S}^{-1}\mathbf{S}_i\mathbf{S}^{-1}\mathbf{S}_k\mathbf{S}^{-1}\mathbf{S}_j + \mathbf{S}^{-1}\mathbf{S}_j\mathbf{S}^{-1}\mathbf{S}_{ik} - \mathbf{S}^{-1}\mathbf{S}_i\mathbf{S}^{-1}\mathbf{S}_j\mathbf{S}^{-1}\mathbf{S}_k\right\} \\ & - I(i=j=k)\exp(b_i), \end{aligned}$$

where $\mathbf{S}_{ik} = \partial^2 \mathbf{S} / \partial b_i \partial b_k$ with jlth element given by

$$\begin{cases} \exp(b_1)\left[\sum_{f=1}^{F}(d_{jf}-d_{lf})^2\right]\left(\exp(b_1)\left[\sum_{f=1}^{F}(d_{jf}-d_{lf})^2\right]-1\right)S_{jl}, & \text{if } i=1 \text{ and } k=1, \\ \exp(b_2)(t_j-t_l)^2\left(\exp(b_2)(t_j-t_l)^2-1\right)S_{jl}, & \text{if } i=2 \text{ and } k=2, \\ \exp(b_1+b_2)\left[\sum_{f=1}^{F}(d_{jf}-d_{lf})^2\right](t_j-t_l)^2 S_{jl}, & \text{if } i=1 \text{ and } k=2. \end{cases}$$

References

References

Conti, S. and O'Hagan, A. (2010), "Bayesian Emulation of Complex Multi-Output and Dynamic Computer Models," *Journal of Statistical Planning and Inference*, 140, 640–651.

Fielding, M., Nott, D., and Liong, S. (2011), "Efficient MCMC schemes for computationally expensive posterior distributions," *Technometrics*, 53, 16–28.

Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., and Rubin, D. (2014), *Bayesian Data Analysis*, vol. 3rd, Chapman and Hall.

Geweke, J. F. (1991), "Effcient Simulation from the Multivariate Normal and Student-t Distributions Subject to Linear Constraints," in *Computer Science and Statistics. Proceedings of the 23rd Symposium on the Interface*.

Girolami, M. and Calderhead, B. (2011), "Riemann manifold Langevin and Hamiltonian Monte Carlo methods (with discussion)," *Journal of the Royal Statistical Society, Series B*, 73, 123–214.

Iserles, A. (2009), *A First Course in the Numerical Analysis of Differential Equations*, Cambridge University Press, 2nd ed.

Johnson, M., Moore, L., and Ylvisaker, D. (1990), "Minimax and maximin distance designs," *Journal of Statistical Planning and Inference*, 26, 131–148.

Kennedy, M. C. and O'Hagan, A. (2001), "Bayesian calibration of computer models (with discussion)," *Journal of the Royal Statistical Society, Series B*, 63, 425–464.

Laidler, K. (1984), "The development of the Arrhenius equation," *Journal of Chemical Education*, 61, 494.

Loeppky, J. L., Sacks, J., and Welch, W. (2009), "Choosing the sample size of a computer experiment: a practical guide," *Technometrics*, 51, 366–376.

Miasojedow, B., Moulines, E., and Vihola, M. (2013), "An Adaptive Parallel Tempering Algorithm," *Journal of Computational and Graphical Statistics*, 22, 649–664.

Overstall, A. and Woods, D. (2013), "A Strategy for Bayesian Inference for Computationally Expensive Models with Application to the Estimation of Stem Cell Properties," *Biometrics*, 69.

— (2015), "Bayesian Design of Experiments using Approximate Coordinate Exchange," *arXiv:1501.00264*.

— (2016), "Multivariate emulation of computer simulators: model selection and diagnostics with application to a humanitarian relief model," *Journal of the Royal Statistical Society Series C*, In press.

Papaspiliopoulos, O., Roberts, G., and Skold, M. (2003), "Non-centered parameterizations for hierarchical models and data augmentation," in *Bayesian Statistics 7*, eds. Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., Oxford.

R Core Team (2015), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

Rasmussen, C. (2003), "Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals," in *Bayesian Statistics 7*, eds. Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., Oxford.

Soetaert, K., Petzoldt, T., and Setzer, R. W. (2010), "Solving Differential Equations in R: Package deSolve," *Journal of Statistical Software*, 33, 1–25.